

A Deep Neural Network and Reconstructed Phase Space Approach to Classifying 12-lead ECGs

David Kaftan¹, Richard J Povinelli²

¹Marquette Energy Analytics, Milwaukee, US

²Marquette University, Milwaukee, US

Abstract

The aim of this work is to classify 12-lead ECGs into 26 classes, including normal sinus rhythm, atrial fibrillation, left bundle branch block, and ST-segment depression and elevation. This work is team Marquette's submission to the PhysioNet/Computing in Cardiology Challenge 2020. Our approach is to apply two modelling techniques: a reconstructed phase space — Gaussian mixture model (RPS-GMM) method and a one-dimensional convolutional neural network. The one-dimensional convolutional neural network consists of 11 layers consisting of both convolutional and fully connected layers. It takes inputs of varying lengths to output a single diagnosis and is trained from scratch within the competition time limits. Our second method, the RPS-GMM approach, embeds each ECG lead into an 11-dimensional space and classifies using a maximum likelihood classifier. While we propose and discuss two methods only the deep convolutional neural network was used in our submissions. The RPS-GMM approach was not scored as it exceeded the competition training time limit. We achieved a score of 0.492 on the test data, but were not ranked due to omissions in the submission. Next steps include reducing the training time of the RPS-GMM approach and ensembling the two methods.

1. Introduction

This work is team Marquette's submission to the PhysioNet/Computing in Cardiology Challenge 2020 [1], [2]. The challenge rules, discussion of the dataset, and challenge results are described in [2]. The rest of our paper focuses on the methods we have used with a short discussion and conclusion section.

2. Methods

We classify 12-lead ECGs into 26 classes, including normal sinus rhythm, atrial fibrillation, left bundle branch block, and ST-segment depression and elevation. We introduce a convolutional neural network (CNN) that is

shallow enough to be trained within the competition time limits and can handle variable input sizes to predict a single diagnosis. A second method used is the reconstructed phase space - Gaussian mixture model (RPS-GMM) approach, which models the ECG signals in a reconstructed phase space using Gaussian mixtures.

2.1. Preprocessing

We perform some initial preprocessing of the signal. First, we make sure the signal is sampled at a consistent rate. We normalize the frequency to 500 Hz using simple down sampling of the 1000 Hz signals that exist in the training set.

We look for extreme outliers by thresholding signals that are 10.5 standard deviations away from the mean. These are replaced by the previous value of the time series. The signals are then z-score normalized for the GMM-RPS method and similarly scaled to 3.5 standard deviations for the CNN.

2.2. Convolutional Neural Network

One-dimensional convolutional neural networks (1D-CNN) are being used to great effect in time series processing [3]. Hannum et al. show a convolutional network can outperform cardiologists in heart arrhythmia detection [4]. They use a single model that takes a 30 second, 200 Hz ECG signal and predicts a single diagnosis for each second of the signal. The method is difficult to reproduce in this competition. It would be difficult to train a model of that size (over 30 layers) in the competition time limit. Also, the method predicts a single diagnosis for every second of the ECG sample; this competition requires multiple diagnoses to be predicted, but the data does not distinguish at what time in the sample the diagnosis occurs. We choose a rather different network architecture to mitigate these problems.




Figure 1: Architecture of the Convolutional Neural Network. Filters and pools are applied to down-sample the original 500Hz signal. Then a global pool is applied to go from an unbound time series to 40 features. The 40 features are used to classify a single diagnosis.

A diagram of our network architecture can be found in Figure 1. A different network is trained for each possible classification. The networks are programmed using TensorFlow [5]. Here, we describe in detail the input, hidden, and output layers.

The input data of the network is a 12-lead ECG signal of variable duration sampled at 500 Hz. Nearly all of the training data is sampled at 500 Hz. Signals sampled at harmonics of 500 Hz are simply resampled by only sampling every i -plus-1-th sample for the i -th harmonic. Signals of any other frequency are resampled using Scipy’s fft resampling method [6]. Signals are crudely cleaned to remove enormous outliers – any sample that deviates from zero by over 10.5 standard deviations is replaced by the sample at the previous time step. Each signal is scaled by 3.5 standard deviations.

The hidden layers of the network all follow the same pattern – apply convolutional layer (rectified linear unit activation, see Figure 1 for size of layers), then down-sample using a max pooling layer (see Figure 1 for size of pools). This pattern is repeated to down-sample the 500 Hz signal to 5 Hz. We also employ a skip connection that simply down-samples the input layer using max pooling without any convolutions. The final max pooling is a global max pool. This allows us to derive 40 features from a variable length time series. We then apply a fully connected hidden layer, then a single class output layer with a sigmoid activation. In summary, several convolutional and max pooling layers “down-sample” the original series.

After the original 12 leads are down-sampled to 5 Hz and 40 features, a global max pool transforms the time-series to a set of 40 features, regardless of time-series length. A fully connected layer is applied along with an output layer. Long short-term memory (LSTM) was tried in place of the global max pool. It did not improve performance and increased training time.

To train each network, we first split our data into training (80% of all data) and validation (20% of all data) sets. Each network is trained on an equal number of positive and negative samples. The Adam optimizer [7] is used with a binomial cross-entropy loss function and a batch size of 1. The validation data is scored every 6800 batches. Training continues until the validation score fails to improve 6 consecutive times. The model with the best validation score is used. We use the geometric mean of F -Beta and G -Beta scores to score our validation data (3). Let TP , FP , and FN be the true positives, false positives, and false negatives, respectively.

$$F_{\beta} = \frac{(1 + \beta^2) * TP}{(1 + \beta^2) * TP + FP + \beta^2 * FN}, \beta = 2 \quad (1)$$

$$G_{\beta} = \frac{TP}{TP + FP + \beta * FN}, \beta = 2 \quad (2)$$

$$Validation\ Score = \sqrt{F_{\beta} * G_{\beta}} \quad (3)$$

2.3. RPS-GMM

Our second method is the RPS-GMM method proposed in [8], [9]. This method captures the natural measure of the attractor generated by the heart. Takens shows that a signal measured from a system can be used to create a state space topologically equivalent to the original state space of that system [10]. The manifold on which the state trajectory lies is modelled statically using a GMM, where a GMM is a set of multidimensional Gaussians whose means and covariances are estimated using expectation maximization [11].

To form the RPS, the dimension of the space, which is the number of samples of the signal, the ECG in this case, is calculated using the global false nearest neighbours method. We use the dimension of 11 learned in [8]. The second parameter is the lag between signal samples, which was learned experimentally on the training set. We use a lag of four. Thus, a lead of the ECG is embedded in an 11-dimensional space with lags of 4 between samples. Let x_n be the n^{th} point in the ECG signal, d be the dimension of the RPS, and τ be the time lag.

$$\mathbf{x}_n = \langle x_{n-(d-1)\tau}, \dots, x_{n-\tau}, x_n \rangle \quad (4)$$

is a point in the reconstructed phase space.




Figure 2. The x-axis is lead 1 of the ECG from record A0001. The y-axis is that signal delayed by 4 samples. The gray points are the points in the RPS. The ellipses are the one standard deviation contours and the cords are in the direction of the eigenvectors of the covariance matrices.

Let $p(\mathbf{x})$ be the probability of a point in the GMM, M be the number of mixtures, w_m be the m^{th} weight, and N be the multidimension Gaussian mixture. Let $\boldsymbol{\mu}$ be the mean of a Gaussian and $\boldsymbol{\Sigma}$ be the corresponding covariance matrix. A GMM is defined as

$$p(\mathbf{x}) = \sum_{m=1}^M w_m N(\mathbf{x}, \boldsymbol{\mu}_m, \boldsymbol{\Sigma}_m) \quad (5)$$

A multivariate RPS is generated using all 12 spaces and a 16 component GMM with full covariance mixtures was learned on each class of ECG signals. An example two dimensional RPS-GMM is shown in Figure 2 for the signal shown in Figure 3.




Figure 3. Z-scored time series of lead 1 from record A0001. The signal is sampled at 500Hz

Test signals are classified by embedding them in the same structured space as the original RPS. The GMM with the maximum likelihood is selected as the class.

3. Results

We achieved a final score of 0.492. We did not officially score due to a late submission of a preprint of this paper. Our methods were otherwise in line with competition rules. These results are for the CNN. The RPS-GMM model entries were not scored because they exceeded the allowed training time limit.

4. Discussion and conclusion

Due to the fast-paced nature of this competition, there are several ideas that we were not able to try. We trained individual models for each classification because we found training a single model with multiple outputs to be unstable. Given the success of other 1D-CNNs in predicting diagnoses using a single model with multiple outputs [4], we expect our problems training could be remedied with more time tweaking hyperparameters. This would have allowed us to create a single larger network while remaining within the training time limit of the competition.

One of the biggest draw backs of our methods is their generic use of the time series. No specific features related to each arrhythmia were used. We feel this could be especially important for ST related heart arrhythmias.

Future work will include addressing the computational complexity for the RPS-GMM approach. This will require modifications to the expectation maximization algorithm. Our last step is to ensemble the RPS-GMM with the CNN.

Address for correspondence:

Richard J. Povinelli
ECE Department, Marquette University
1515 W. Wisconsin Ave.
Milwaukee, WI 53233
US
richard.povinelli@marquette.edu

References

- [1] Ary L. Goldberger *et al.*, “PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals,” *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000.
- [2] E. A. P. Alday *et al.*, “Classification of 12-lead ECGs: the PhysioNet/Computing in Cardiology Challenge 2020,” *Physiol. Meas.*
- [3] S. Kiranyaz, O. Avci, O. Abdeljaber, T. Ince, M. Gabbouj, and D. Inman, “1D convolutional neural networks and applications: A survey,” *ArXiv*, vol. abs/1905.0, 2019.
- [4] A. Y. Hannun *et al.*, “Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network,” *Nat. Med.*, vol. 25, no. 1, pp. 65–69, 2019, doi: 10.1038/s41591-018-0268-3.
- [5] Martin Abadi *et al.*, “TensorFlow: Large-scale machine learning on heterogeneous systems.” 2016, [Online]. Available: <https://arxiv.org/abs/1603.04467v2>.
- [6] P. Virtanen *et al.*, “SciPy 1.0: Fundamental algorithms for scientific computing in python,” *Nat. Methods*, vol. 17, pp. 261–272, 2020, doi: <https://doi.org/10.1038/s41592-019-0686-2>.
- [7] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Proceedings of 3rd International Conference for Learning Representations*, 2015.
- [8] R. J. Povinelli, M. T. Johnson, A. C. Lindgren, F. M. Roberts, and J. Ye, “Statistical models of reconstructed phase spaces for signal classification,” *IEEE Trans. on Signal Process*, vol. 54, no., pp. 2178–2186, 2006, doi: 10.1109/TSP.2006.873479.
- [9] R. J. Povinelli, M. T. Johnson, A. C. Lindgren, and J. Ye, “Time series classification using Gaussian mixture models of reconstructed phase spaces,” *IEEE Trans. Knowl. Data Eng.*, vol. 16, no. 6, pp. 779–783, 2004, doi: 10.1109/TKDE.2004.17.
- [10] F. Takens, “Detecting strange attractors in turbulence,” in *Dynamical Systems and Turbulence*, 1980, vol. 898, pp. 366–381.
- [11] B. S. Everitt and D. J. Hand, *Finite Mixture Distributions*. London ; New York: Chapman and Hall, 1981.